

AI STRATEGY · FRAMEWORK

Productising AI Decision Framework

For product and leadership teams deciding where, if anywhere, AI belongs in the product, packaging, and pricing.

How to use this: Work through each section for any proposed AI feature before committing engineering time, then use the go/no-go summary to make the call explicit.

The Customer Job

Start from the job to be done, not the technology.

What specific customer job or pain point is this AI capability meant to address?

What is the customer doing today without it, and how painful is that really?

What evidence do we have that customers would value this, beyond assumption?

Where This AI Capability Belongs

Identify which of the four placements fits, since each has different implications.

Placement type	Fits this use case? (Y / N)	Implication if chosen
Discrete feature		
Packaging / tier differentiator		
Pricing input (usage-based cost)		
Experience layer (no explicit 'AI' label)		

Reliability Bar

What reliability level does this feature need to hit before it's customer-ready, and why?

What happens, visibly, when the AI gets it wrong, and is that an acceptable customer experience?

What honest language will we use to set expectations without overselling certainty?

Cost and Pricing Implications

- ☐ Model/usage cost of this feature has been estimated at realistic customer volumes
- ☐ A decision has been made on whether this cost is absorbed, tiered, or usage-priced
- ☐ Margin impact at scale has been reviewed, not just at pilot volume

Go / No-Go Decision

Decision: proceed, proceed in a narrower form, or not yet, and why?

If proceeding, who owns this from here through launch?

If not yet, what needs to be true before we revisit this?
